

Claims

What is claimed is:

1. 1. In a Constant Access Time Bounded (CATB) cache, a method comprising:
 2. reserving a first number of unallocated lines in the cache for pinned data, the first number being less than the number of lines in the cache; and
 3. if data needs to be inserted into the cache as pinned data,
 4. selecting a line from the lines reserved for pinned data;
 5. storing the data in the line; and
 6. inserting the line into a search group of the CATB cache.
1. 2. The method of claim 1 wherein each line of the cache is stored in non-volatile memory.
1. 3. The method of claim 2 further comprising:
 2. recovering the organization of the cache on power up following a loss of power to the cache by
 3. in a first phase of recovery, for each line in the cache
 4. determining if the line is allocated;
 5. if the line is allocated, inserting the line in a search group of the cache; and
 6. if the line is not allocated, inserting the line into a pool of free lines;
 7. and
 8. in a second phase of recovery, for each search group
 9. determining the number of pinned lines in the search group; and

11 adding at least one line from the pool of free lines to each search group
12 that has at least one pinned line.

1 4. The method of claim 3 wherein the cache is a disk cache in a processor based
2 system.

1 5. The method of claim 1 wherein inserting the line into a search group of the
2 cache further comprises:

3 indicating that the line is allocated;
4 indicating that the line is pinned; and
5 using a tag of the line to map the line to a search group of the cache.

1 6. The method of claim 5 wherein:

2 the CATB cache is implemented as a set-associative cache;
3 each search group of the cache is a set of the cache; and
4 inserting the line into a search group of the cache further comprises:
5 using the address of the data as the tag of the line;
6 performing a modulus operation between the tag and the number of sets
7 (N) in the cache (the tag MOD N) to map the tag to a set of the cache;
8 performing a search based on the tag of the line; and
9 inserting the line into a dynamic data structure that represents the set.

1 7. The method of claim 6 wherein indicating that the line is pinned further
2 comprises modifying metadata associated with the line to indicate that the line is
3 pinned.

1 8. For a whole number N, in an N-way set associative non-volatile disk cache, a
2 method comprising:
3 reserving a predetermined number of lines for pinned data and organizing them
4 into a pool of lines for pinned data;
5 distributing the remaining lines in the cache into N dynamic data structures of
6 approximately the same size to represent the N sets of the cache;
7 if data is to be inserted into the cache as pinned data,
8 inserting the data into a line from the pool for pinned data;
9 marking the line as allocated by modifying metadata associated with the
10 line;
11 determining the set to which the line belongs using a mapping based on
12 the tag associated with the line;
13 removing the line from the pool for pinned data; and
14 adding the line to the set.

1 9. The method of claim 8 further comprising:
2 recovering the organization of the cache on power up following a loss of power to
3 the cache by
4 in a first phase of recovery, for each line in the cache
5 determining if the line is allocated;
6 if the line is allocated, inserting the line in a set of the cache using a
7 mapping based on the tag associated with the line; and

8 if the line is not allocated, inserting the line into a pool of unallocated
9 lines; and
10 in a second phase of recovery, for each set in the cache
11 determining the number of pinned lines in the set using the metadata
12 associated with each line in the set; and
13 moving one or more lines from the pool of unallocated lines to each set
14 that has at least one pinned line so that the number of non-pinned lines in
15 each set is approximately the same.

1 10. An apparatus comprising:

2 an N-way set associative cache implemented in non-volatile memory

3 a pinned data portion of the non-volatile memory to store a pool of lines for

4 pinned data; and

5 a pinned data insertion module to

6 insert pinned data into a line from the pool of lines for pinned data;

7 mark the line as being allocated by modifying metadata associated with

8 the line;

9 determine a set to which the line belongs using a mapping based on the tag

10 associated with the line;

11 remove the line from the pool for pinned data; and

12 add the line to the set.

1 11. The apparatus of claim 10 further comprising
2 a power source to provide power to the cache; and

3 a recovery module to recover the organization of the cache on power up following
4 a loss of power to the cache from the power source by
5 in a first phase of recovery, for each line in the cache
6 determining if the line is allocated;
7 if the line is allocated, inserting the line in a set of the cache using a
8 mapping based on the tag associated with the line; and
9 if the line is not allocated, inserting the line into a pool of unallocated
10 lines; and
11 in a second phase of recovery, for each set in the cache
12 determining the number of pinned lines in the set using the metadata
13 associated with each line in the set; and
14 moving one or more lines from the pool of unallocated lines to each set
15 that has at least one pinned line so that the number of non-pinned lines in
16 each set is approximately the same.

1 12. A system comprising
2 a processor;
3 a disk communicatively coupled to the processor;
4 an N-way set associative cache implemented in non-volatile battery-backed up
5 Dynamic Random Access Memory communicatively coupled to the processor;
6 a pinned data portion of the non-volatile flash memory to store a pool of lines for
7 pinned data; and
8 a pinned data insertion module to

9 insert pinned data into a line from the pool of lines for pinned data;
10 mark the line as being allocated by modifying metadata associated with
11 the line;
12 determine a set into which the line using a mapping based on the tag
13 associated with the line;
14 remove the line from the pool for pinned data; and
15 add the line to the set.

1 13. A machine readable medium having stored thereon data which when accessed by
2 a machine causes the machine to perform the method of claim 1.

1 14. The machine readable medium of claim 13 having stored thereon further data
2 which when accessed by a machine causes the machine to perform the method of
3 claim 2.

1 15. The machine readable medium of claim 14 having stored thereon further data
2 which when accessed by a machine causes the machine to perform the method of
3 claim 3.

1 16. The machine readable medium of claim 15 having stored thereon further data
2 which when accessed by a machine causes the machine to perform the method of
3 claim 4.

1 17. The machine readable medium of claim 13 having stored thereon further data
2 which when accessed by a machine causes the machine to perform the method of
3 claim 5.

- 1 18. The machine readable medium of claim 17 having stored thereon further data
- 2 which when accessed by a machine causes the machine to perform the method of
- 3 claim 6.

- 1 19. The machine readable medium of claim 18 having stored thereon further data
- 2 which when accessed by a machine causes the machine to perform the method of
- 3 claim 7.

- 1 20. A machine readable medium having stored thereon data which when accessed by
- 2 a machine causes the machine to perform the method of claim 8.

- 1 21. The machine readable medium of claim 20 having stored thereon further data
- 2 which when accessed by a machine causes the machine to perform the method of
- 3 claim 9.

- 1 22. In a Constant Access Time Bounded (CATB) cache, a method comprising:
 - 2 initializing a search group of the CATB cache with a capability to dynamically
 - 3 insert and delete elements; and
 - 4 inserting elements dynamically into the search group of the CATB.

- 1 23. The method of claim 22 further comprising:
 - 2 receiving a first identifier for an element;
 - 3 using the first identifier to compute a second identifier for a search group in the
 - 4 CATB cache; and
 - 5 traversing the search group to locate an element matching the first identifier.

6 24. The method of claim 23 wherein the search group is implemented as a linked list.

1 25. A machine readable medium having stored thereon data which when accessed by
2 a machine causes the machine to perform the method of claim 22.

1 26. The machine readable medium of claim 25 having stored thereon further data
2 which when accessed by a machine causes the machine to perform the method of
3 claim 23.

1 27. The machine readable medium of claim 25 having stored thereon further data
2 which when accessed by a machine causes the machine to perform the method of
3 claim 24.